Chapter 5

RETHINKING TIME AND DETERMINISM: WHAT HAPPENS TO DETERMINISM WHEN YOU TAKE RELATIVITY SERIOUSLY

Jenann Ismael

Columbia University

Ji2266@columbia.edu

Determinism is a centrally important notion for physics: it links time to laws and connects events along spatial surfaces to events along the temporal dimension. In the context of space-time theories, failures of determinism have been viewed as pathologies and used to identify superfluous structure. In philosophy, determinism has played its most important role in discussions of free will, where a certain picture of what determinism entails has a strong grip on the imagination. According to that picture, a deterministic universe unfolds with physical necessity from an initial condition that was set long ago. This presents a strong challenge to our sense of agency because it takes two very basic commitments — the idea that the laws of physics place fundamental constraints on what can happen (you throw a ball in the air or set a pendulum in motion and you know exactly what is going to happen) and that the past is fixed — and it uses the laws to leverage the fixity of the past into the fixity of the future. Neither of those commitments seems negotiable. There's a famous argument that makes this explicit that goes, in simple terms, like this: the past is fixed and out of our control; the laws are fixed and out of our control.

Our actions are entailed by those together, so they are fixed and out of our control [1].

Physics has changed since Newton's time. The changes I will be interested in come from relativity and not from quantum mechanics. Since we are likely too big, too warm, and too slow for quantum effects to show up in our behavior, classical physics is likely to be the effective physics for understanding the human being. Relativity makes a difference because determinism is a relationship between time and laws, and relativity transforms our conception of both.[1] This is a chapter about what happens to determinism when one takes relativity seriously.

Relativity introduces a number of changes in our conception of time. Past and future become relativized to a point. Instead of a division of events into past, present, and future, the Universe is divided at each point into causal past, absolute elsewhere, and causal future. The logic of these notions is different from the logic of pre-relativistic notions of past, present, and future in significant ways. Temporal development no longer makes sense at the level of the Universe because there is no longer a global notion of "now." It still makes sense to speak of temporal development along the worldline of an embedded observer and from that perspective, things remain largely unchanged. The change relativity introduces in our conception of laws is that the laws are formulated in local terms. Everything is the product of local events and interactions. An event A can influence an event B only if they can be connected by a process that propagates at finite speed through the intervening space. Although the laws remain deterministic in the sense that the total state of the Universe at a time determines its total state at any other (or rather — translated into relativistically well-defined vocabulary — every global spacelike hypersurface is a Cauchy surface), the causal past at a point no longer determines events even a finite fraction of a second into the future. I will argue that the upshot of these changes is that although the laws remain deterministic, relativity eliminates — at the level of the geometry — the point of view from which the fixity of the past can be leveraged into the fixity of the future.

This is not a chapter intended to solve the problem of free will. It is a chapter about temporality in a relativistic world. The point of starting with the discussion of the problem of free will was to bring into focus the

[1]Newtonian physics contains an ambiguity relevant to the interpretation of determinism that relativity resolves by giving us a way of formulating laws that are purely local.

way of thinking of determinism that I want to dismantle. It's bearing on the discussion of free will is that it provides a step towards greater clarity.

1 What I Will Do in What Follows

I will develop these points and their philosophical interpretation more slowly in what follows. After introducing determinism and the special threat it has been taken to pose to human freedom, I'll show how relativity forces us to alter pre-theoretic assumptions about the difference between space and time, the nature of temporal development, and the form of the physical laws. We will see how these changes challenge the common understanding of determinism.[2]

Fatalism: The problem of fatalism was around long before relativity, but it gained affirmation in some people's minds from the relativistic image of time. Classical fatalism, associated with Parmenides, is a statement about the nature of time. We have a pre-theoretic view of the world as a spatially extended thing with a fixed past and an open future of branching possibilities. The claim is that relativity treats time just like the spatial dimensions and that sense of openness is illusory. Past, present, and future all exist already, and there's no good sense in which they remain to be fixed.

Determinism: Determinism, by contrast, has to do with physical laws. Again, there were some early versions of the problem raised by seeing human action as under the sway of physical laws but it really received a precise and pressing form in Newtonian mechanics, where it became possible to write down equations that would take a description of the state of the Universe as a whole at one time and calculate its state at any other.

Fatalism and determinism seem like different problems. The former concerns time, specifically with a conception of time associated with relativistic theories. The latter doesn't present itself as having anything particularly to do with time and certainly not the specific view of time that relativity introduces. Indeed, it's usually discussed in an entirely

[2] I speak almost entirely about Minkowski spacetime, the spacetime of special relativity. This is the simplest setting and the one in which the points come through most clearly. The general relativistic regime adds complications without changing the substance. See also note 12.

pre-relativistic setting. It has to do with laws, which are thought to provide a necessary connection between the state of the world at different times.

The idea behind the worry that seems raised by determinism is that if you start with the idea that the past is fixed but the future is open, determinism would leverage the fixity of the past into the fixity of the future. This is the idea that I want to examine.[3]

2 Does Relativity Commit you to Fatalism?

First, a few remarks about whether the 4D image of time that relativity asks you to adopt commits you to fatalism. When one makes the geometry of space-time part of the proper subject matter of physics, one has to semantically ascend to a form of representation that takes all of space and time — including past, present and future — into its representational scope. To give imaginative content to what the theory is saying, we tend to lapse into metaphors that embed low-dimensional representations of space-time in time, and then imagine occupying a godlike perspective outside looking down at it. People are almost irresistibly inclined to say things like that on a 4D representation the future is already there, that it already exists, and so the idea that there is any contingency is a mistake.

The response to all of this is to point out that the 4D representation is just what history looks like when time is made endogenous. It is a harmless kind of representational transformation that doesn't commit us to thinking that the future is *already* there. In a 4D representation, things happen when they happen and at no other time, just as they do in a 3D representation, and just as they do in life. Nor is the transcendent perspective literally a point of view from which your future could *now* be visible to an appropriately situated observer. There is no external temporal dimension and hence no point outside the manifold from which one could even formally construct a point of view. And if the *now* in that

[3] I leave the assumption that the past is fixed unexamined here. The philosophical challenge to free will usually takes the common sense idea that the past is fixed for granted and uses it (together with determinism) to challenge the fixity of the future. In physics, however, the question of whether and in what sense the past is fixed can no longer be assumed uncritically. The combination of relativity and the time symmetry of the classical mechanical laws have made it something that demands physical explication. The bulk of current scientific opinion favors a thermodynamic conception. See [2–5].

phrase ("from which your future could now be visible to an appropriately situated observer") refers to the internal time of our world, the claim is simply false. Your future is visible from the future for beings in the future, and nobody else. Those are all just confusions that come from metaphors that we use to try to give imaginative content to the transcendent vision by creating a low-dimensional representation of time *in time* and then looking down at it and judging things from the perspective of the embedding time.

The transcendent perspective on time that relativistic theories ask us to adopt does not have deep philosophical consequences for free will. It is committed to the fatalistic statement "what will be, will be" only in the tautologous sense that it is ***not*** the case that what will be, will ***not*** be. It is most definitely not committed to the idea that the future is *already* there, or that there are facts about the way the world is *already*, in the here and now, that necessitate that things will be as they will be. That is why determinism seems to pose a much stronger and more worrisome problem. Precisely because determinism does seem to entail that there are facts in place already — things that have happened already, not in the future, but in the here and now — that necessitate that things will be as they will be and it's the physical laws that supply the necessity.

3 What Determinism does (and doesn't) Entail

So, is it the case that in a deterministic world there are facts in place already in the here and now that necessitate events in the future? When people introduce both these problems they often ask us to consider what is true at different times. When they introduce fatalism, they often say things like "It is true today that there will be a sea-battle tomorrow, and so it can't be possible today that there isn't a sea-battle tomorrow." When they introduce determinism they often say that it was true at the beginning of time that the laws were such and such and the initial conditions were such and such and since everything else follows logically from the laws and initial conditions, the future was fixed from the beginning of time. In both cases the talk of what is true at different times invites confusion. What people are really trying to capture with that vocabulary are questions about what is fixed already or fixed by facts that are already in place, but the vocabulary of truth is simply not suited to do that work. It makes no distinction between first-order facts and higher-order facts, between intrinsic and relational facts, between facts that are local to here and now and facts that describe happens at other times and places. I'm going to avoid all talk about what is true at different times. I'll speak only in terms of the first-order facts that form the subject matter of all of this talk: what happens here (at some moment) and what happens there, and what the physical laws entail about the relationship between what happens here and what happens there. Or, to put it into relativistic language, I'm going to talk about nomological (or lawful) relationships between the intrinsic contents of one volume of space-time and the intrinsic contents of another. I'll sometimes call this the "on-the-ground causal order," but this should be understood as simply shorthand for what the laws entail about the relationships between the contents of different volumes of space-time. [4]

Is it the case that the intrinsic contents of the past in a deterministic world determine, as a matter of physical law, what happens a minute, a day, a year after that? We know that in a deterministic theory, there are equations that let us calculate the state of the Universe at one time, given its state at another. So, for example, in Newtonian mechanics, if we know the positions and momenta of all the particles of which the Universe is composed (determinants of its total state), we can calculate their positions and momenta at every other time. But it turns out that determinism does not entail that there are facts already in place at any given moment that nomologically (i.e., as a matter of law) determine or necessitate what happens at future times. More precisely, determinism does not entail that the microscopic past, characterized in intrinsic, first-order terms, nomologically determines the future.

That is easy to show. Consider a deterministic, Newtonian universe. Consider the collection of events that happen at some early time, and an event e that occurs some time later. Indeed, consider the whole history of the world up until t, and an event that occurs just a fraction of a second after t. Do the events that make up the history of the world up until t — lets call the collection of those events C — nomologically determine e? That is the question of whether there are solutions to the Newtonian equations of motion which include C and *not* e? The answer is yes. To obtain

[4] "Causal" is being understood here in the stripped-down sense that is connected to geometry: two events are causally connectible if they can be connected by a physical process. A causal interaction is a local interaction that involves the exchange of a conserved quantity and a physical process is a continuous chain of interactions. See [6]. Reference [7] gives an overview of the different notions that are available in physics and the connections among them.

such a solution we just add things to C that change the forces impinging on e. It's not the specific collection of events in C, but the totality of events in the Universe as a whole that determines e. There's no nomological necessity linking the events in C to the total state. That is quite general. For any collection of pre-t events, there are models of the Newtonian equations of motion that include those events but lead to very different futures. In order for a set of past events to nomologically determine a future event, we have to specify that those events constitute the totality of what there is.

That works in a Newtonian space-time because there is no restriction on past events that might be relevant to some future event. For reasons related to this, Minkowski space-time is usually regarded as the most hospitable environment for determinism. *The problem in Newtonian space-time is that there was no way of explicitly specifying (in first-order terms) the collection of events that is sufficient to nomologically (as a matter of law) determine some future event. In Minkowski space-time, there is.* The only spatiotemporal structure is the light cone structure which separates time into three regions at any point p: the absolute past (which contains the collection of events that can be connected to p by a light signal, and hence the collection of events that can be causally relevant to p), the absolute future (which contains the collection of events that can be causally affected by what happens at p), and the absolute elsewhere (which contains the collection of events that cannot effect or be effected by what happens at p; there is no signal or influence that connects events at p to events in its absolute elsewhere).

It turns out that the situation for determinism is no better in that setting when we consider events that lie at even a tiny finite interval in the future.[5] It *is* true in Minkowski space-time that any given event e is nomologically determined by the events in an explicitly specifiable volume of space-time (its back-light cone, or any temporal cross section of it). But if we consider an event at any finite interval in the future, the causal past of the latter includes events that are *not* in the causal past of the earlier. If you look at the two events t_0 and t, here you can see that the causal past

[5] In a Newtonian setting, precisely because there is no finite speed of propagation, there is no way of explicitly delimiting a set of events that are nomologically sufficient to produce some event e. This is remedied in Minkowski space-time where the causal past of an event contains the full set of events that that can have any causal bearing on its occurrence. Minkowski spacetime gives a more explicit rendering of the causal order and allows us to unambiguously connect spatiotemporal order to causal order.

of the latter includes events that are *not* in the causal past of the earlier. The light cones are nested inside one another so that there is space between them, and this space includes events that are nomologically necessary for the later event and not included in the past of the former.

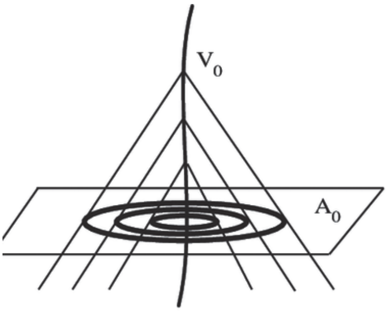

This diagram shows the back-light cones of events along the worldline W projected onto the spacelike hypersurface A_0. This is the relativistic realization of the common-sense idea that tomorrow, of course, has a different causal past than today, because it includes all of the events that have happened in between. The difference is that "in between" has a new relativistic reading. It includes events that fall outside the back-light cone of yesterday, but that common sense would think of as in the distant past. In Minkowski space-time looking *pro*spectively, the laws themselves tell us that there are indefinitely many ways of extending the causal past. Looking *retro*spectively, any event is fixed by its own causal past. The information that becomes available in that interim period between the prospective and retrospective points of view (information about what happens between now and then) is not itself constrained by the causal past and is essential to nomological fixation of the future.

There are two (nomologically interchangeable) ways of characterizing what happens in between. One can take an early cross section of the back-light cone and evolve it forward, or one can just give the initial state and the exogenous variables. In both cases, there are new events that aren't part of the causal past and that are nomologically necessary to fixing the future. If p_0 is here-now, p is an event in the absolute future of p_0, and {e*} is the set of events that fill out the space between the two light cones, then for any of the events in {e*}, there are frames in which it occurs before p_0 and frames in which it occurs after, and as a class the events in {e*} aren't all fixed until p.

4 Temporal Development in a Relativistic Setting

The only meaningful order internal to the Universe is the causal order embodied in the light cone structure. There is a well-defined order for events that fall in one's past causal horizon and a well-defined order for events that fall in one's future light causal horizon, but no well-defined order (relative to the present moment) for events that fall in the absolute elsewhere. Because there is no notion of "now" that spans the whole universe (and correspondingly, no unequivocal notion of past and no such notion of future), there is no temporal development at the level of the Universe as a whole. Past and future are relativized to a point and temporal development is relativized to a worldline. Temporal development along a worldline is obtained by comparing the perspective from different points along that line. So you compare how the past and future look from t_n, t_{n+1}, $t_{n+2...}$ In a pre-relativistic world, temporal development is characterized by a kind of expansion in the future direction. In a relativistic world, temporal development is relativized to a worldline of an observer is a kind of expansion upward and outward. The past expands along the spatial dimensions as we move up the temporal dimension (i.e., as we compare perspectives at later moments along a timelike curve). As one gets more future, one also gets more past.

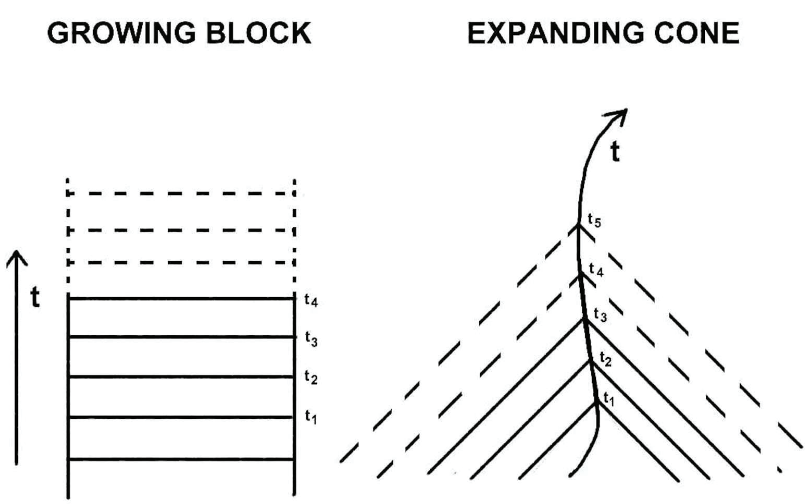

Relativity breaks down the intuitive difference between time and space, bringing spatiotemporal structure more explicitly into line with causal structure (past and future are defined by relations of causal connectibility) and eliminating the point of view from which the fixity of the past can be leveraged into the fixity of the future. When past and future are reinterpreted in a relativistic setting, determinism no longer entails that that past determines the future. To be explicit about what the reinterpretation of past and future in a relativistic setting involves: it involves relativizing to a point in space-time, identifying "the past" with the past light cone (the causal past) and "the future" with the future light cone (the causal future). The motivation for the reinterpretation is that these are the classes of points from which influence can reach p, and from which p's influence can be felt, respectively. *Points in the absolute elsewhere have no temporal order relative to p — they happen neither before nor after — precisely because they have no causal order relative to p*. Points in the absolute elsewhere have no direct physical relationship to p or its past; they are unconnected in the causal network defined by p and its past. They become connected in the future when information from them crosses into p's light cone. They become connected in p's future when the causal network that includes p expands to include them as well.

The lesson that points in the absolute elsewhere have no temporal order relative to p is not an incidental feature of the relativistic space-time structure. It stems from the fundamental innovation of the theory. Time isn't an external parameter in which the Universe unfolds. It is one of the dimensions in which the Universe is extended, along with the three spatial dimensions, and the intrinsic geometry of the four-dimensional object doesn't give us a well-defined notion of the state of the Universe at a time. You can talk about past and future at a point. But just as there is no univocal notion of now, there is no univocal notion of past and no univocal notion of future. The past light cones of spatially distant points do not coincide. The past of someone in your absolute elsewhere will include points not in yours, and vice versa. The time it takes for a signal to travel from there to you is the time it takes for your pasts to coincide. And there is no physically meaningful sense in which the events in the past are *already* in place on their way to meeting you at the crossroads. If you are using temporal vocabulary, you have to relativize it to a point; "already" means "in the past," and these events are not "already" from your point of view, in place.

5 The Logic of Past and Future in a Relativistic Setting

There is a thing one can do, however, which I signaled already above in talking about the relationship between the prospective and retrospective views in a relativistic world and which is worth dwelling on because it is jarring to the pre-relativistic imagination. Choose any inextensible timelike curve (any timelike curve that is as extended as it could possibly be), go to the last point along that curve and then take a cross section of the back-light cone of that last moment shortly after the Big Bang. That cross section would capture the causal past of all events that would ever make an impact on points along the worldline. In an infinite universe, there would be no last moment, but we could extend the worldline to future infinity, and we take a cross section of the back-light cone of that point. So described, it is tempting to think that even if it comes fully into view only at the end of time, what is revealed at the end is — surely — what things were like *already* at the beginning. There may have been a basis for saying that in Newtonian space-time, but in a relativistic universe, there is absolutely no basis for thinking that events in our absolute elsewhere are "already there" and on their way to cross into our future light cone. The events in the absolute elsewhere that are (as a class) nomologically sufficient to fix a future event e happen no more before e than after it, and are no more or less fixed than e itself. They don't, in any objective sense, *happen before* e. The thing that gives meaning to temporal order is the causal order, and because these events have no causal bearing on e, they have no temporal order relative to it.

One can also describe this in a local way: looking back from my 80th birthday, the retrospective perspective from that point — obtained by taking a cross section of my past light cone that includes my birth — will contain events sufficient to nomologically determine all of the exogenous influences that would impact me over the course of my life. In that sense, even though my past at the time of my birth doesn't nomologically determine its future, my future *post-determines* a past that does. So when you reach the end of your life, you can look back and it will look to you as if it couldn't have gone any other way — even though it was all novelty, i.e., leaps of blind faith, and openness looking forward — given your starting point and all of the things that happened to you along the way. It will look to you as though what you end up becoming is what it was always destined to be. In fact, however, the only sense in which that thought is true is the fatalistic one. The events that look to you retrospectively as though they were *always* going to happen the way that they did are — in no good sense — *already* in place before you act. There is nothing inconsistent here. That is just part of the logic of past and future in a relativistic universe. It reinforces why one can't rely on pretheoretical intuitions to draw inferences in this setting.

The way to think of things is that every event in a relativistic universe is the confrontation of lines of development leading from widely separated points along the spatial dimension. And the order of those confrontations is dictated by the geometry. The history of the Universe *is* the history of these meetings — Einstein called them "point coincidences." There is no meaningful sense in which time itself passes and no meaningful notion of temporal development at the level of the world as a whole. Every life unfolds one event at a time with a past that never predetermines its future, but a future that post-determines its own past. All of this makes perfect sense and the reason that makes sense is because of the way the light cones are next inside one another.

6 The Universe as a Causal Network

Pre-theoretically, there is a large intuitive difference we draw between space and time. If you cast your eyes across a landscape, for example, you think of the landscape as a fixed object that simply comes into view in stages. We tend to think of time, by contrast, as coming into being as it is experienced. And when I describe this picture of development as growth along both dimensions, you instinctively think of the growth along the spatial dimensions as simply revealing what was "already there." Determinism plays on this pre-relativistic distinction between space and time. The idea was that if you thought the future was open in a way that the past was not, determinism was supposed to force you to give that up by showing that the future is nomologically determined by the past. The move from Newtonian to relativistic physics eliminates the pre-theoretic sense of a distinction between space and time.[6] There is a shift from an

[6] Instead of thinking of the Universe is not a spatially extended object evolving in time, but a four-dimensional structure of events. The difference between space and time comes to be understood in terms of the more subtle differences between the spacelike and timelike dimensions (the minus sign attaching to timelike dimension in the signature of the metric).

object-based to an event-based ontology. The basic entities are events. Geometry now encodes relations of causal connectibility. There remain differences between the spacelike and timelike dimensions, but the basis for treating space as a substance and time as the dimension of becoming is gone. The new setting also lets us draw the distinction between what is *fixed* and what is *open* in a more refined way. Where before "fixed" meant something like "exists already," in the new setting, "fixed" means something like "lies in the causal past, i.e., effects, but is not effected by, what happens in the here-now."[7] "Not yet fixed" means "hinges on what happens here and now." The result of these changes is the picture that I described above of growth along both dimensions, with events in the absolute elsewhere having exactly the same status as events in the future, and both inheriting a kind of openness from their connection to action. It is often said that relativity spatializes time. It might better be said that it temporalizes space. We can no longer think of the distant parts of space as "there already," but coming gradually into view. The spatial past is just as infected by "becoming" as the temporal future.

Let me repeat this in a slightly different way. The information about exogenous variables that is needed to get from p_0 to p_1 is contained in the world in two nomologically interchangeable ways: it is contained in the tails between the two back-light cones of p_0 and p_1 if we go back to a moment after the Big Bang (the earliest moment at which the deterministic equations become applicable),[8] or is distributed across the worldline connecting p_0 to p_1. This also goes if we let p_0 be the initial and p_1 be the final state of the Universe. What this means is that the information that is contained along a cross section of the back-light cone of this last moment is not contained in the past of any earlier, but rather spread along the whole temporal dimension between the initial and final moments. In this way, information about the total state of the Universe along the spatial dimensions, judged from any point in space-time, is always information *from the future*: information from the causal future, which is the only notion of future that is physically meaningful here. It is not that the future, judged from some moment p, is *pre-determined* by p's past. It is that

[7] This intuitive idea is something that demands physical explication and will acquire a much richer meaning when combined with a thermodynamic gradient. It is not something that can be understood purely in geometric terms, so I leave it unanalyzed here. See [3] and [8].

[8] More generally, if we draw a spacelike hypersurface at any time that cross-sects both light cones, it is contained in the events along that surface between the two light cones.

points in p's future *post-determine p's* causal past. That and *more*, in fact. The retrospective perspective from the last point on a given worldline will not in general capture the causal past of the Universe as a whole, because the back-light cones of spacelike related points are not coextensive.[9] In general, we don't get a truly *global* view unless we consider the retrospective view from that last moment along *all* inextensible timelike curves in the Universe. And that in turn means that *there is no global view of space-time from within*.

The most natural way to think of a relativistic world is as a causal network. Basic interactions are interactions that involve the local exchange of a conserved quantity. Physical processes are chains of basic interactions and events are causally connectible if they can be connected by a physical process. Events are temporally locatable with respect to one another when they fall in one another's causal future or causal past and there is a well-defined temporal order for events along any given worldline, but a slight asynchrony between events on worldlines distant from one another. There is no global time. The intrinsic causal structure of the world is a network of worldlines with a little bit of give that leaves an asynchrony between worldlines at a spatial distance. Each point in the network has a view of its own causal past. No point in the network has a view of the whole. The distinctive worry that determinism seemed to present — that we could use the physical laws to leverage the fixity of the past into the fixity of the future — is eliminated by the geometry. Geometry no longer provides a point of view from which the fixity of the past can be leveraged into the fixity of the future. Past and future are not related as slices through the world that separate the spatial past from the spatial future, but as stacks of nested light cones that piece together into a four-dimensional Minkowskian structure. The causal past of points in that structure does not determine their future. The only point of view that spans all of space also spans all of time. Relativity rules out, at the level of the very geometry, a point of view that seems to allow determinism to pose the distinctive worry.

It cannot be an accident that when relativity brought spatiotemporal structure more into line with causal structure, it removed the point of view from which the fixity of the past can be leveraged into the fixity of

[9] The exception is general relativistic space-times that have a so-called God Point, i.e., a point in whose causal past the whole space-time falls. See [9].

the future. Minkowski space-time effectively creates a way station between the past and future (viz. the absolute elsewhere) for events that have not occurred yet, have no temporal order relative to here and now, are no more past than they are future, and are nomologically essential to determining what happens next. These events aren't *there* for us yet, in that they are not connected in the causal network defined by our past. They come into our world when information from them crosses into our future light cone.

7 Laws and Conservation Principles

To this point, we've been talking mostly about geometry, but there are associated developments in the form of physical laws. In relativistic theory, physics becomes local. The theory's laws relate what happens at a point only to its immediate environment. An event in one part of space-time can affect an event in another only by a physical process that passes through the space between. Physical processes are themselves chains of basic interactions, each of which involves the exchange of a conserved quantity. No process propagates at a speed faster than the speed of light. This is a clear advance over Newtonian mechanics because it eliminates one of the primary obscurities of the theory. In Newtonian mechanics, the motion of a particle depends on the total force acting on it. Since the law of gravity says that every particle attracts every other particle in the Universe with a force that is directly proportional to the product of their masses and inversely proportional to the square of the distance between them, to know the force on any given particle, one would have to know the position and momentum of every other particle in the Universe. People objected to this feature of the theory almost immediately when it was published. It wasn't simply philosophical resistance to the idea of action at a distance; no mechanism was provided by which a particle at the other end of the Universe exercised influence on a particle here and Newton famously declined to comment ("Hypotheses non fingo"). In a relativistic Universe, everything that happens in the universe is the product of local events and interactions. The processes that underlie the nomological relationships between particles in different parts of space and time are openly rendered.

The second important development — related to the one above — is that the status of global conservation principles get clarified. Conservation laws have always played an important role in physics. The natural way to think of the conservation of energy, for example, pre-relativistically, is there is a fact about the total energy of the Universe that is fixed at some initial time and that constrains the global energy going forward. The same goes for momentum, angular momentum, electric charge, or any conserved quantities. That certainly mirrors the way that one calculates: one writes down something that represents the total energy, and the laws evolve it forward. According to this line of thought, the total energy and other conserved quantities have an absolute value that was fixed at the beginning of time and constrains how the Universe as a whole evolves. The view of time that underwrites that interpretation is no longer tenable in a relativistic theory. Time is now conceived as one of the internal dimensions of the Universe and the intrinsic geometry of the four-dimensional manifold doesn't give us a well-defined notion of the total energy at a time.[10] There were other reasons why the status of conservation principles became a serious concern when Einstein introduced the equations of the general theory of relativity. This was sorted out when Emmy Noether proved a very beautiful theorem that changed how we think of conservation laws. Noether proved that every differentiable symmetry of the action of a physical system has a corresponding conservation law. This suggests that conservation laws are not irreducibly global constraints, but rather artifacts of local dynamical symmetries. If the local exchanges of energy (for example) are conservative — i.e., if you never raise the energy of an open system without taking it from somewhere else, or lower it without depositing it somewhere else — energy will be globally conserved for the Universe as a whole simply because there is no place outside from which to draw it.

This gives us a clue about how determinism can hold globally while failing from the perspective of any embedded system inside the universe. So long as the laws that govern the basic interactions entail that there is no difference between the future of two systems that have the same (causal) pasts unless there is a difference in the influences impinging from outside, then the universe as a whole will be deterministic, even though the global state wasn't fixed in advance, and even though nothing on the

[10] One might be tempted here to cite the fact that the total energy is the same for any foliation. While that's true, it is not clear what the significance this fact has outside of a Newtonian conception of time. It doesn't support a physical interpretation of global states in the past constraining the state moving forward and raises its own mystery of why — in the absence of that interpretation — the total state on any foliation should be conserved. In either case, one needs to understand conservation in terms of local processes.

ground is "keeping track" of the global state (so to speak). That's simply because the Universe is the causal network itself. By definition, there is nothing that is not included in it, so there are no exogenous influences to draw on to generate differences in the future. That in turn means that any kind of internal freedom that is compatible with the local laws will have global determinism as a corollary. Global determinism adds nothing to local constraints.

So, locally, the past does not determine the future. To the extent that global determinism does hold, it turns out to be an artifact of local symmetries. Looked at from an internal perspective — i.e., from an on-the-ground point of view in space and time — novelty and contingency is the rule. Each system in the world is developing in response to a barrage of influences, new information arises at every stage of its development: external contingencies that come from nowhere and can't be foreseen. Every system is channeling energy and transforming exogenous influences into behavior. That goes for drainage systems and drosophila as surely as for penguins and persons. When I say that those external contingencies come from nowhere, I mean, they come from the absolute elsewhere. If we are viewing the world at point p, those events are outside the relational network in which p is connected. Those events will become *locatable* with respect to p in p's future when information from them crosses into its light cone. And any tendency to think that they were "lying in wait all along" is to be resisted because those events occur no more before e than they occur after. There is nothing in the history of a point that determines what will happen to it tomorrow, a year from now, or a century. Relativistic physics supports the common-sense idea that it all depends on what happens in between. We could get together with our friends and family, fixing our causal pasts and everything that we could in principle know and bring together in one time and place, and it remains (from a nomological point of view) open what is coming down the pike. You and I, and frogs and flowers, even mountains and oceans, are open systems, each with our own causal past, responding to local influences in a world that none of us has a global view of. We collectively create a future that isn't fixed by any of our own causal pasts.

8 Taking Stock

So, going back to the question of whether determinism entails that the past determines the future. Is it the case that past events in a deterministic world nomologically determine future events? The answer is "no." Strictly speaking, it's not true in Newtonian space-time, and it's not true in Minkowski space-time, which is generally thought to be the most hospitable setting for determinism, because there is a limit that lets us explicitly proscribe the set of events that can be causally relevant to some future event. In Newtonian mechanics, we saw that it wasn't true because, although the total state of the world at a time nomologically determines its state at other times, the total state wasn't fixed by any particular collection of events. For any given past, we could always find multiple solutions to the Newtonian equations with wildly different futures by embedding them in a larger world. The situation was somewhat equivocal in that setting, however, because the total state of the Universe at any initial time does determine everything that happens thereafter, so we have an ambiguity: the total state in the past determines the future but not the intrinsic contents of the past.[11] Minkowski space-time clarified the situation by removing the ambiguity. The laws make no reference to total states; they provide a clear and explicit rendering of the on-the-ground causal order (as well as the causal processes that underlie all nomological relationships). The set of events nomologically sufficient to determine a particular event is now explicitly proscribed and it is unequivocally true in that regime that determinism does not entail that the causal past at any point determines an event that falls even a tiny fraction of an interval in the causal future. The situation in general relativity is more complicated. In general relativity, the local structure is still Minkowskian, so all of the points about local causal structure carry over. The curvature, however, is no longer generally Minkowskian. The fundamental laws of the theory — Einstein's field equations — relate the curvature to the matter content, and this leads to a wide variety of physically possible global topological structures, among them space-times that are not temporally orientable, space-times that contain

[11] The question of when and by what the total state is fixed is more difficult than one would think. In logical terms what is needed to fix the total state, given a collection of events, is the negative fact that there are no others, or what is sometimes called a "totality fact": i.e., the claim that the facts accounted for constitute the totality of what there is. The question then becomes when and by what *these* kinds of facts are determined. The status of these kinds of facts in a setting like the one here in which there are no intrinsic maxima on the size of the world is philosophically contested because one can find no first-order supervenience base of local matters of particular fact. This is why the elimination of reference to totality from the physical laws and reformulation in local terms in relativity clarifies the situation. For more on totality facts, see for example [10–12].

closed timelike curves, and space-times that contain holes and other pathologies. These kinds of topologies open up the space for other failures of determinism.[12] Because of this, general relativity has never been a good home for determinism.

9 The Shift in Vision

Time is an external parameter in our lives, individually and collectively. And it is an external parameter in the histories of the systems around us. We watch rocks and rivers, bees and barn animals, timing their movements using our own internal clocks. When we scale up to the level of the Universe, we do the natural thing, treating it as a big object and time as an external parameter in its evolution. And later when we learn about determinism, it seems to entail that everything is fixed by the fundamental laws together with the initial conditions of the Universe. It seems that once the initial conditions are established, the Universe unfolds with strict necessity. It seems that the chaos and unpredictability of life as we experience is a product of incomplete information and the prohibitive complexity of calculating.

Relativity teaches us that the natural way of scaling up was mistaken. Time isn't outside the Universe. It is *inside*. The Universe is a spatiotemporally extended manifold ordered by relations of causal connectibility and defined by a light cone structure. A relativistic world is not a big object unfolding in time but a collection of ongoing processes unfolding asynchronously relative to one another. If we trace a path along the worldline of any physical object, what happens in the future will not

[12] See [13]. One might worry that the existence of finite models of general relativity jeopardizes the claim that the past doesn't determine the future. I don't think it does. While it is true that if one fixes the global topological structure and the events in the causal past of a point that includes a Cauchy surface, then one fixes the future nomologically, the crucial question is whether the global topological structure is determined by the causal past of any given point, and the answer is no. This follows from a theorem proved by John Manchak, following a conjecture of David Malament [14], which shows that the causal pasts of all points in any model of general relativity can be embedded in another model with a different global topological structure. So it remains true that the causal past of any event is compatible with indefinitely many nomologically possible futures. The point is related to the issue about the total state in Newtonian physics. The status of these global properties (when and by what they are fixed) is one that hasn't received enough attention.

be determined by the causal past. Events will come seemingly from nowhere, bringing information from their own causal past. The future will be jointly determined by the intersection of the pasts of all of the points that meet at a given point, and the order of these meetings will be determined by the geometry. The determinist's belief that one could go back far enough into one's own past and cast one's net widely enough to capture initial conditions rich enough to predict the evolution of the world as a whole is simply mistaken.

10 Recap, and Brief Remarks About Free Will

To sum up the innovations that relativity introduces and make some remarks about their bearing on the debate about free will: In Minkowski space-time, notions of past and future become relativized to points, replacing the pre-relativistic division of events into past, present, and future with the division into causal past, absolute elsewhere. This is not a conservative change: the notions have a different logic than that of past, present, and future. The idea of temporal development at the level of the Universe does not survive. There is a meaningful notion of temporal development along the worldline of an observer, obtained by comparing the perspective from different points along that line. From that perspective, temporal development is much as common sense thinks. One's life unfolds one event at a time, and at every moment there's a clean division of the events of one's life into past, present, and future. The innovation is an asynchrony between temporal development along worldlines of distant observers. Fundamental laws previously formulated as global laws of temporal development are reformulated in local terms. Conservation laws and global principles like determinism turn out to be artefacts of local processes. The result of these changes leads to the elimination of the point of view from which the fixity of the past can be leveraged into the fixity of the future.

This lesson is masked by practices in physics. Formally, of course, everybody recognizes that there is no physically preferred foliation. It is still common practice in physics, however, to suppose that this just gives us a little leeway in how we carve up the Universe and *modulo* freedom in choice of a foliation, we can speak sensibly about the Universe's development in time. We still speak routinely of the initial conditions and identify

Cauchy surfaces in the distant past. We represent the Universe by pairing a Hamiltonian that represents its total energy with an operator that gives the law of evolution. When we want to introduce a possible universe, we often fix the total state along spatial dimensions and use the laws to generate the future. These practices paper over the deep innovation that the theory introduces and that is directly relevant to the interpretation of determinism. Taking relativity seriously means recognizing that that mixed point of view — where we fix the total state along the spatial dimensions in the past and try to leverage that into the fixity of the future — is not well defined. There is no such point of view. The only point of view that spans all of *space* also spans all of *time*. There is unequivocally no global present, no initial moment at which all the events nomologically sufficient to determine the future were put in place,[13] and the total state of the world is not in general fixed by the causal past of any point in space-time. Once we eliminate reference to the total state, the actual, immanent causal structure internal to space-time emerges. What we find is that in a deterministic setting in a relativistic space-time is that events in the set needed to determine an event e are not fixed prior to e, but exactly at e.

How does this affect the discussion of free will? For the purposes of understanding free will, this was a ground-clearing exercise. It undermines the simplest and most intractable version of the challenge that determinism presents to free will.

The simple version of the challenge is the one usually presented in popular discussions or in philosophy classes. It is the one people tend to encounter first, and it is so powerfully persuasive because it is so simple. It starts from a single premise that is firmly rooted in common sense — that the past is fixed — and can seem to make any detailed discussion of the place of human choice in the on-the-ground causal order is moot/ beside the point/ a merely epistemic kind of openness/ at best a consolation for the lack of the deeper kind of openness and at worst wretched subterfuge that fails to address the real problem. It is, I will confess, the

[13] In the physics literature, Lee Smolin [5] has been the most powerful advocate for preserving originality, spontaneity, creativity, the openness of the future (all of the things that he thinks of as essential to the "reality" of time), and he thinks that to do that you have to reject both determinism the central lesson of relativity and cling to a preferred foliation. I'm suggesting the opposite: that really appreciating the relativistic rejection of a preferred foliation eliminates the external perspective from which we can talk objectively about the total state of the Universe at a time.

version of the argument that most bothered me. No matter how sophisticated one gets in defining freedom or analyzing what it means to "have the ability to do otherwise than you in fact do," if everything that happens to you is destined by facts that are in place long before you came into the world, it is hard to see how any of it could really address the problem.

The most abstract way to put the simple version of the challenge that determinism presented is that there is an asymmetry between past and future — expressed by the idea that the past is fixed and the future is open — that is essential to the agential perspective. The idea that the past is fixed is most naturally glossed as the conviction that nothing that we can do here and now can change the past, and it is underwritten imaginatively by a metaphysical picture of a universe that is unfolding in time. Add determinism to this picture and it makes it impossible to maintain the asymmetry between past and future. One need not be specific about what "fixity" amounts to; determinism in a Newtonian setting it will force you to say that in whatever sense the past is fixed, so is the future. It will force you to give the future the same status as the past. I put this earlier by saying that determinism allows one to leverage the fixity of the past into the fixity of the future.

The simple version of the challenge is defeated by the considerations that we have been discussing.[14] In a relativistic setting, past and future are

[14] Those steeped in philosophical literature will want to know how to relate the so-called Consequence Argument, which has been the central argument for the incompatibility of free will and determine since the 1980's. What I am calling the Simple Version is what I take to be the intuitive heart of the argument, but the argument has been given numerous formulations and the literature surrounding it is too tangled and arcane to engage usefully here, so I'll just make a few remarks. One starts from the presumption that for an agent to be acting freely in her performance of an action it must be possible for her to do otherwise than she does in fact. The Consequence Argument is supposed to show that no agent in a deterministic world has the possibility of doing otherwise. A lot hinges on how one formulizes the ability to do otherwise. Early versions of the argument formalized the requirement on freedom in ways that make it vulnerable to solutions that trade on the logic of counterfactuals, or exploit an ambiguity in how the past is characterized. (See [15], [16], and [18].)

The strongest version of the argument, and the one that gets closest to capturing the initial intuition, is given in (17) and it formulates the requirement thus:

(PFPL) An agent *S* has it within his power in a possible world *w* to do *X* at time *t* only if there is a possible world with the same (temporally intrinsic) past relative to *t* and the same laws as in *w* in which *S* does *X* at *t*.

The considerations in this chapter address this version of the argument, and show that this requirement for freedom is satisfied in a relativistic world. The causal past of a point p relativized to points, temporal relations among points separate into three classes (past, future, and absolute elsewhere), the laws are rendered locally, and the past of a point no longer determines its future. The metaphysical picture of a universe unfolding in time, moreover, is no longer viable. The situation is completely different. The effect is to open up a space in which freedom might breathe.

It remains true that when an agent is looking at some action A that lies in her future at a point p, the events in her past light cone at p, *together with exogenous influences that will impinge on her between p and A*, will fix what happens at A. And one might be tempted to say, "The problem was never really about past and future; it was really about control. Our actions are entailed by things out of our control: "the past" was just a stand-in for events out of our control. Here's a new argument: the future of an agent at a moment m is determined by the past light cone of m plus exogenous variables that impinge on her over the course of her history. Neither of these is under *her control, so her future is not under her control*. But when we look more carefully at this new argument, things are not so clear. We can no longer find a nomologically sufficient basis in the distant past from which to derive the future, so we are forced to look at the embedded agent and her role in the production of behavior. The agent herself is now recognized as an embodied system, a space-occupying object whose own activities are part of the causal order. To assess the argument, we need some physical model of the agent, so let us model her schematically as a minded body whose voluntary movements are controlled by the activity in the brain, noting that the activity we are particularly concerned with is the activity that supports conscious deliberation and decision making.[15] If we ask whether the agent's behavior at A is determined by the exogenous variables impinging on her body between now and then, the answer is no. One has to take into account the state of her brain and body at the initial moment, and if we schematize what is going on in her body between p and A, with special attention to the tole that the brain processes underwriting deliberation and decision are playing, it will emerge that those processes are regulating the impact of exogenous influences on voluntary

at which an agent faces the choice about a future action is compatible with physically possible worlds in which an agent performs indefinitely many different actions.

[15] We assume that conscious mental supervenes on activity in the brain (but is not necessarily identical with it) because this is the assumption that brings human action under the scope of physical laws raising the threat posed by determinism.

behavior. From an information-processing perspective, deliberation is a process in which the thoughts, memories, beliefs, hopes and dreams, values, priorities, and so on that are encoded in her brain are brought to bear in deciding what to do, so the physics here actually vindicates the common-sense view that one's decisions control the bearing of exogenous influences on one's actions. None of this should be controversial from a physical point of view. If one is treating the agent as a physical system and assessing the role that the brain processes underwriting deliberation and decision play in the production of human behavior, these conclusions are supported by what we know.[16] As soon as we are forced to explicitly represent the agent, it becomes clear that an agent's deliberations and decisions are integral to the future development of the world and that they play a special role in regulating the bearing of exogenous influence on behavior.

Attention will then turn to the fact that the state of the agent's brain, and all the information encoded in the state of the brain at any moment, is itself the product of its past. And the worry, if there is one, will have to be that in a deterministic setting, the process in which those were extracted from what happens to the agent over the course of her life was itself deterministic. Addressing this concern will take some care and the kind of care that it demands will raise all of the deepest questions about what we are and what it means for an action to come from us. It will confront questions about the role we play in constituting ourselves out of the noisy accidents of our lives and force us to reflect on how much of what we are is the product of our own choices and how much is the product of biology and circumstance. It will engage with questions about what we can be held responsible for, what it means to be the kind of thing that can take responsibility, what it means to be autonomous and reflective, and many others besides. Thinking through these issues will force us to confront difficult philosophical questions and to make choices about what freedom really is.

This is what I meant by saying that this chapter was a ground-clearing exercise. It pushes the discussion of free will out of the realm of physics and into the subtleties and complexities of philosophy, which is where I think it belongs. So long as the simple version of the argument from determinism remained in place, there was (not for everybody, but for some of us) a sneaking suspicion that the philosophical discussion was

[16] See [8] for a fuller discussion.

largely moot.[17] Now it looks, I believe, absolutely to the point. It's all that there is.

11 Conclusion

This chapter has been about what happens to determinism when one take relativity seriously. The formal observations are easily verified. They are just points of physics. The philosophical interpretation that I have suggested should accompany them is a thorough rejection of the half-way internal/half-way external point of view that fixes the total state along spatial dimensions and tries to think of the Universe as an object evolving in time. The remnants of that way of thinking — one that makes space endogenous and time exogenous, i.e., the one that treats time as an external dimension in which the universe evolves instead of as one of the dimensions in which the universe is extended — are still everywhere in physics and philosophy, but the greatest philosophical contribution of the theory — the one that we should try to take to heart — is that there is no such point of view.

References

[1] Van Inwagen, Peter, 1975. The incompatibility of free will and determinism. *Philosophical Studies*. 1975 27(3): 185–199.

[2] Albert, David Z. *Time and Chance*. Cambridge, MA: Harvard University Press; 2000.

[3] Albert, David Z. *After Physics*. Cambridge, MA: Harvard University Press; 2015.

[4] Maudlin, Tim. Remarks on the passing of time. *Proceedings of the Aristotelian Society*, New Series 2002; 102. p. 259–274.

[5] Smolin, Lee. *Time Reborn*. Boston, MA: Houghton Mifflin Harcourt; 2013.

[6] Dowe, Phil. Causal process theories. In: Beebee H., Hitchcock C., Menzies P. (eds.), *The Oxford Handbook of Causation*. Oxford, UK: Oxford University Press. 2012.

[7] Ismael, Jenann. Causation: An overview of our emerging understanding. In Kleinberg, Samantha (ed.), *Time and Causation in the Sciences*. Cambridge, UK: Cambridge University Press; 2019.

[17] I am immensely grateful to Carlo Rovelli for some very helpful exchanges and probing questioning.

[8] Ismael, Jenann. *How Physics Makes Us Free*. New York, NY: Oxford University Press; 2016.
[9] Manchak, John. *Global Spacetime Structure*. Cambridge, UK: Cambridge University Press; 2020.
[10] Armstrong, David. *A World of States of Affairs*. Cambridge, UK: Cambridge University Press; 1997.
[11] Rosen, Gideon. Metaphysical dependence: Grounding and reduction. In B. Hale & A. Hoffman (eds.), *Modality: Metaphysics, logic, and epistemology*. Oxford, UK: Oxford University Press. 2010. p. 109–136.
[12] Russell, Bertrand. *The Philosophy of Logical Atomism*. Chicago, IL: Open Court. 1985. p. 1–125.
[13] Doboszewski, Juliusz. Epistemic holes and determinism in classical general relativity. *British Journal for the Philosophy of Science*. 2020; 71(3):1093–1111.
[14] Manchak, John. Can we know the global structure of spacetime? *Studies in History and Philosophy of Modern Physics*. 2009; 40: p. 53–56.
[15] Whittle, A. Dispositional abilities. *Philosophers Imprint*. 2010; 10: 1–23.
[16] Fischer, J. M. *The Metaphysics of Free Will*. Oxford, UK: Blackwell Publishers; 1994.
[17] Fischer, J. M. Local-Miracle Compatibilism: A Critique. In: Hausmann M., Noller J. (eds.), *Free Will*. Cham: Palgrave Macmillan. 2021.
[18] Fischer, J. M., and Pendergraft, G. Does the consequence argument beg the question? *Philosophical Studies*. 2013; 166 (3): p. 575–595 (2013).